**≥DSU**

AD-A246 882

||||||||||||||||||||||

FINAL REPORT
Schemas in Problem Solving:
An Integrated Model of Learning, Memory, and Instruction

Sandra P. Marshall

December 1991

DTIC
ELECTE
MAR 0 4 1992
S B D

Center for Research in Mathematics and Science Education

College of Sciences
San Diego State University
San Diego, CA 92182-0413

92-05117

||||||||||||||||||||||

REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorates for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	January 1992	Final Report 10/1/89 - 12/31/91

4. TITLE AND SUBTITLE	5. FUNDING NUMBERS
FINAL REPORT Schemas in Problem Solving: An Integrated Model of Learning, Memory, and Instruction	N00014-89-J-1143 G 442c010--6

6. AUTHOR(S)

Sandra P. Marshall

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
Department of Psychology/CRMSE San Diego State University San Diego, CA 92182	CRMSE Report 91-02

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSORING/MONITORING AGENCY REPORT NUMBER
Office of Naval Research Cognitive Science Program (Code 1142CS) 800 N. Quincy Street Arlington, VA 22217-5000	

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT	12b. DISTRIBUTION CODE
Approved for public release; distribution unlimited.	

13. ABSTRACT (Maximum 200 words)

This final report contains two papers. The first describes statistical and cognitive models used to simulate student performance. The statistical model provides information about how the group of students as a whole performed on an identification task involving word-problem situations and shows differences among subgroups. The cognitive model is a connectionist network that simulates the performance of each student and yields details about how learning varied from one individual to another. The second paper outlines a hybrid model of schema knowledge that joins a connectionist network with a production system. Details of the model are provided, and an example of its output is presented.

14. SUBJECT TERMS			15. NUMBER OF PAGES
cognitive model, problem solving, hybrid model			61
			16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	UL

FINAL TECHNICAL REPORT

Grant No: ONR N00014-89-J-1143
Period: October 1, 1989 - December 31, 1991
Date of Submission: January 25, 1992
Name of Institution: San Diego State University
Title of Project: Schemas in Problem Solving: An Integrated Model of
                         Memory, Learning, and Instruction
Principal Investigator: Sandra P. Marshall

Table of Contents

Project Summary

This document serves as the final technical report of ONR Grant No. N00014-90-J-1143, which was funded as a continuation of work carried out under ONR Contract No. N00014-85-K- 0661. As part of the initial project I developed a schema-based model of teaching and learning for the domain of arithmetic word problems (Marshall, Pribe, & Smith, 1987). The schemas emphasize the basic situations that can be contained in such problems. A central focus of the research was to create a model that applied equally well to issues of memory organization, teaching and learning, instructional development, and diagnosis of student learning.

A core set of situations was identified, and a series of studies verified that the situations were sufficient for describing virtually all legitimate word problems (Marshall, 1990). A model of schema knowledge was constructed for each of the basic situations. Each schema model specified the feature knowledge, constraint knowledge, planning knowledge, and implementation knowledge required to use the schema successfully. An extension of the basic schema model yielded ways in which affective components may also be part of schema knowledge (Marshall, 1989). Attention was also given to ways in which different types of schema knowledge could be easily assessed (Marshall, 1988). More recently, I have demonstrated that the schema theory can be applied easily to two other domains, elementary statistics and rational number instruction (Marshall, in press a).

The instructional system, called *STORY PROBLEM SOLVER (SPS)*, was designed to provide instruction about these situations in such a way as to foster the development of appropriate schemas by individuals (Marshall, Barthuli, Brewer, & Rose, 1989). The system consists of (a) a series of lessons requiring about 6-8 hours for completion and (b) a flexible problem-solving environment. Both of these components were designed to focus on specific aspects of schema knowledge required in solving problems.

In the lessons, each component of schema knowledge was addressed implicitly through short instructional segments and related exercises. Students were introduced to a set of icons depicting the situations, and they were encouraged to use the icons to represent the various situations occurring in specific problems. A set of experiments revealed that students did develop the specific types of schema knowledge targeted by SPS and that the icons were a key part of their knowledge (Marshall & Brewer, 1990). Moreover, we were able to chart the development of schemas over the course of instruction through individual interviews with our subjects (Marshall, in press b).

The second part of the system is a flexible Problem-Solving Environment, *PSE*, in which students can experiment with problem representations by manipulating the icons described above. Students are able to select a subset of icons to represent a problem and to link these together to represent the connections in the problem. They have options to expand the icons and explore individual aspects of each one, to carry out calculations, to select other icons if they so desire, or to have the system display a possible representation of the problem. This environment was developed under the

original ONR Contract and evaluated under the project continuation as Grant N00014-90-J-1143 (Marshall, 1991).

The project yielded three major products. First, I created a working computer-based system of instruction that can be used to teach students about solving word problems. The system has been used successfully with about 100 subjects to date (primarily college students with weak problem solving skills). Second, I have developed and refined a theory of schema structure and acquisition. The theory builds on the general nature of schema knowledge found in the cognitive science and cognitive psychological literature but goes considerably beyond it. In particular, the theory allows operational definition of key components of a schema and thus allows empirical tests of whether individuals have acquired these pieces. Third, as a direct consequence of studying the acquisition of schema knowledge and attempting to evaluate students' learning, I have formulated a new model of assessment. The model is a network model, and it stipulates the need for assessing both the number of nodes and the connectivity within the net. Thus, the project results allow us to use the theory of memory organization (i.e., schema theory) to model learning, instruction, *and* assessment. This last result has had the most far-reaching impact. As can be seen from the attached list of publications and presentations, I have been invited to make a number of contributions about assessing schema knowledge. The importance here is that the theory developed during this project is unique in its use of a common model for learning, instruction, and assessment. Moreover, the theory provides the basis for a linkage between a psychological theory of memory/learning and a new psychometric theory of testing.

Finally, the project also yielded several important modeling *results*. We have simulated successfully the performance of students as they respond to the computer exercises. The simulation uses estimates of their schema knowledge as revealed in interviews. Both correct and incorrect responses are equally well estimated. We have also employed a series of connectionist models which learn to classify the situations expressed in story problems. The modeling continued under the project renewal and is the focus of the report which follows. The report has two sections. The first section describes statistical and cognitive models of performance on the initial recognition task in the instructional system. The models described therein successfully simulate actual student performance on an item-by-item basis. The second section describes the full model of schema instantiation. This is a hybrid model, incorporating both a production-system and a connectionist network. It successfully evaluates multi-step story problems, recognizes their important relational components, and solves the problems for the correct numerical solution.

At the end of this summary is a list of publications, technical reports, conference presentations, and invited addresses that report research from these two projects. Much of the work spanned both of them. Most of the research results will be reported in a book now being prepared for publication by the Cambridge University Press. The book should be completed by September 1992.

Project Publications and Reports

Publications:

Marshall, Sandra P. (1988). Assessing problem solving: A short-term remedy and a long-term solution. In R. I. Charles & E. A. Silver (Eds.), *The Teaching and Assessing of Mathematical Problem Solving.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Marshall, Sandra P. (1989). Affect in schema knowledge: Source and impact. In D. B. McLeod & V. M. Adams (Eds.), *Affect and mathematical problem solving.* New York: Springer Verlag.

Marshall, Sandra P. (1990). The assessment of schema knowledge for arithmetic story problems: A cognitive science perspective. In G. Kulm (Ed.), *Assessing higher order thinking in mathematics.* Washington, D.C.: AAAS. [a]

Marshall, Sandra P. (199.)). Generating good items for diagnostic tests. In N. Frederiksen, R. Glaser, A. Lesgold, & M. Shafto (Eds.), *Diagnostic Monitoring of Skill and Knowledge Acquisition.* Hillsdale, NJ: Lawrence Erlbaum Associates. [b]

Marshall, Sandra P. (in press). Assessing schema knowledge. In N. Frederiksen, R. Mislevy, & I. Bejar (Eds.), *Test Theory for a New Generation of Tests.* Hillsdale, NJ: Lawrence Erlbaum Associates. [a]

Marshall, Sandra P. (in press). Assessment of rational number understanding: A schema-based approach. In T. Carpenter, E. Fennema, & T. Romberg, *Rational Numbers: An Integration of Research.* Hillsdale, NJ: Lawrence Erlbaum Associates. [b]

Marshall, Sandra P. (in press). Statistical and cognitive models of learning through instruction. To appear in Meyrowitz, A. L. & Chipman, S. (Eds.), *Cognitive Models of Complex Learning.* Norwell, MA: Kluwer Academic Publishers. [c]

Technical Reports:

Marshall, Sandra P., Pribe, Christopher A., & Smith, Julie D. (1987). *Schema Knowledge Structures for Representing and Understanding Arithmetic Story Problems.*

Marshall, Sandra P. (1988). *Assessing Schema Knowledge.*

Marshall, Sandra P. (1988). *Schema Knowledge for Solving Arithmetic Story Problems: Some Affective Components.*

Marshall, Sandra P., Barthuli, Kathryn E., Brewer, Margaret A., & Rose, Frederic E. (1989). *STORY PROBLEM SOLVER: A schema-based system of instruction.*

Marshall, Sandra P. (1991). *Computer-Based Assessment of Schema Knowledge in a Flexible Problem-Solving Environment.*

3

Presentations:

Marshall, Sandra P. (1987, April). Knowledge representation and errors of problem solving: Identifying misconceptions. In W. Montague (Chair), *Diagnosing Errors in Science and Mathematics*. Symposium conducted at the Annual Meeting of the American Educational Research Association, Washington, D.C.

Marshall, Sandra P. (1988, April). Assessing schema knowledge. In N. Frederiksen (Chair), *Test Theory for Tests Based on Cognitive Theory*. Symposium conducted at the Annual Meeting of the American Educational Research Association, New Orleans.

Marshall, Sandra P. (1989, January). The assessment of schema knowledge for arithmetic story problems. In G. Kulm (Chair), *Perspectives and Emerging Approaches for Assessing Higher Order Thinking in Mathematics*. Symposium conducted at the Annual Meeting of the American Association for the Advancement of Science (AAAS), San Francisco.

Marshall, Sandra P. (1990, April). What students learn (and remember) from word problem instruction. In S. Chipman (Chair), *Penetrating to the Mathematical Structure of Word Problems*. Symposium conducted at the Annual Meeting of the American Educational Research Association, Boston.

Marshall, Sandra P. (1991, April). Computer-based assessment of schema knowledge in a flexible problem-solving environment. In H. F. O'Neil, Jr. (Chair), *Extending the Frontiers of Alternative Assessment with Computer Technology*. Symposium conducted at the Annual Meeting of the American Educational Research Association, Boston.

Invited Addresses:

"Remedial Instruction for Arithmetic Story Problems: A Cognitive Science Approach." Invited address to the National Council of Teachers of Mathematics, Chicago, April 1988.

"Schema Knowledge". Invited presentation to the Resource Center for Science and Engineering, University of Puerto Rico, November 1989.

Unpublished Manuscripts:

Marshall, Sandra P. (1991). *Understanding the situations of arithmetic word problems: A basis for schema knowledge.*

Marshall, Sandra P. & Brewer, Margaret A. (1990). *Learning from icons: What you see is what you get, or is it?*

Book in Preparation:

Working title: *Schemas in Problem Solving: An Integrated Model of Instruction, Learning, and Assessment.* To be published by Cambridge University Press.

Statistical and Cognitive Models of
Learning through Instruction[1]

Sandra P. Marshall
Department of Psychology
San Diego State University
San Diego, CA 92182-0315

Abstract

This chapter uses statistical and cognitive models to evaluate the learning of a set of concepts about arithmetic word problems by a group of students. The statistical model provides information about how the group of students as a whole performed on an identification task involving word-problem situations and shows differences among subgroups. The cognitive model simulates the performance of each student and yields details about how learning varied from one individual to another. It is a connectionist model in which the middle layer of units is specified *a priori* for each student, according to the student's level of understanding expressed in an interview. The chapter concludes with a detailed comparison of the simulated responses with the observed student responses.

INTRODUCTION

The learning investigated here occurred as part of a study in which students received computer-based instruction about arithmetic word problems. The central topics of the instruction were five basic situations that occur with great frequency in word problems: *Change, Group, Compare, Restate,* and *Vary*. The instruction had three main segments: (1) the introduction, in which the situations were described; (2) an in-depth exploration, in which details of each situation were elaborated and presented diagrammatically; and (3) the synthesis, in which combinations of situations

---

were introduced together with planning and goal-setting techniques.[2] For each of the three parts of instruction, students engaged in multiple practice exercises. The study reported in this chapter concerns only the first segment of instruction--the introduction to the situations--and the primary focus is the nature of the knowledge that individuals gained from that introductory instruction.

This chapter describes two analyses of what individuals learn from instruction. Both analyses are needed. In the first case, learning is examined in a traditional experimental paradigm, using established statistical procedures. Group features, rather than individual characteristics, receive greater emphasis in this paradigm, and conclusions drawn from the analysis describe group commonalities. In the second case, learning is examined by means of a cognitive model that simulates individual performance. In this analysis, individuals' characteristics are studied, and conclusions apply separately to each individual. As I indicate below, the information gained from each analysis is valuable in a study of learning. Neither one alone provides the complete picture.

The questions of interest in the research are *what* is the new knowledge retained in memory as a result of instruction, *when* is it retained, and *which parts* of it are later accessed and retrieved. During instruction, some new information is (presumably) acquired and added to an individual's available knowledge store. Not all possible information is taken in, and individuals vary in the type and amount of new knowledge that enter memory. It is the rare instance in which all learners learn exactly the same thing from a single instructional lesson. More often, some learners noticeably remember a great deal of the new information while others remember almost nothing.

The Instructional Domain

This section provides a short description of the five situations used in instruction. The situations are *Change, Group, Compare, Restate,* and *Vary,* and they represent uniquely almost all simple stories found in arithmetic story problems (Marshall, 1991).

The *Change* situation is characterized by a permanent alteration over time in a measurable quantity of a single, specified thing. Only the quantity associated with one thing is involved in the *Change* situation. It has a

---

[2] Details about the computer-based instruction can be found in Marshall, Barthuli, Brewer, & Rose (1989), a technical report available from the author.

beginning state and an end state, with some intervention which causes a transition from beginning to end. Usually, three numbers are of importance: the amount prior to the change, the extent of the change, and the resulting amount after the change has occurred.

A *Group* situation is present if a number of small distinct sets are combined meaningfully into one large aggregate. Thus, the *Group* situation reflects class inclusion. The grouping may be explicit or implicit. If explicit, the solver is told in the problem statement which small groups are to be united. If implicit, the solver must rely on his or her prior semantic knowledge to understand the group structure. For example, in a situation involving boys and girls, the solver would typically be expected to know that boys and girls form a larger class called children. The solver also would be expected to understand that the members of the subgroups (i.e., boys or girls) retain their identity even when combined into a larger group (i.e., children). Three or more numbers are necessary in a *Group* situation: the number of members in each of the subgroups as well as the overall number in the combination.

The *Compare* situation is one in which two things are contrasted to determine which is greater or smaller. The numerical size of the difference between the values is unimportant and may not even need to be computed. The *Compare* situation relies heavily on prior knowledge that individuals have about relations. Most frequently, the *Compare* situation requires the solver to choose either the larger or smaller of two values when the operative relation is stated as a comparative adjective or adverb (e.g., faster, cheaper, shorter, more quickly). The objective is the determination of whether one's response should be the larger or the smaller of the known values. This situation most typically occurs as the final part of a multi-step item. For instance, one often sees problems in which the solver is expected to decide after several problem-solving steps which of two items offered for sale is the better buy. This final determination is a *Compare*. It requires only the recognition of which of the two items is less costly—it does not require the computation of how much less.[3] Most *Compare* items involve values for only two objects, although it is certainly possible to make comparisons among three or more.

---

[3] It should be noted that the *Compare* situation defined here differs from the semantic relation of the same name developed by Riley, Greeno, and Heller (1983).

The *Restate* situation contains a specific relationship between two different things at a given point in time. The relationship exists only for the particular time frame of the story and cannot be generalized to a broader context. There are two determining features of a *Restate* situation. First, the two things must be linked by a relational statement (e.g.,one of them is twice as great as, three more than, or one half of the size of the other).Second, the relationship must be true for both the original verbal descriptions of the two things and the numerical values associated with them. Thus, if Mary is now twice as old as Alice, then 20 years--which is Mary's age--must be twice as great as 10 years, which is Alice's age. Note that this relationship was not true one year ago nor will it necessarily be true in five years.

The *Vary* situation is characterized by a fixed relationship between two things that persists over time. The two things may be two different objects (e.g., boys and girls) such that one can describe a ratio as "for every boy who could perform x, there were 2 girls who could do the same ....", or they may be one object and a measurable attribute of it (e.g., apples and their cost) with the problem having the form "if one apple cost $.50 then five apples ....". An essential feature of the *Vary* situation is the unchanging nature of the relationship. If one of the objects is varied, the amount of the second changes systematically as a function of the known relationship. The variation may be direct or indirect.

Simple examples of these five situations are given in Table 1. During the entire course of computer instruction, each of the situations is introduced, explained, and transformed to a problem setting. Eventually, several are linked together to form multi-step problems. In the introductory lesson, each situation is described by means of an example and with the general features which define it.

Although they are very simple and readily understandable, the five situations are not intuitively known by students through previous instruction. Experiments with groups from several different student populations indicated that students (and teachers) do not typically recognize or use situational knowledge in story problems (Marshall, 1991). Those same experiments show that students of all ages are nevertheless able to learn them.

The present study was designed to investigate how that learning comes about. Because they were previously unknown to the students, the situations in story problems were, in fact, five new concepts to be learned. Thus, the study described here provides a setting for investigating how

Table 1
The Five Situations

CHANGE	To print his computer job, Jeffrey needed special paper. He loaded 300 sheets of paper into the paper bin of the laser printer and ran his job. When he was done, there were 35 sheets of paper left.
GROUP	The Psychology Department has a large faculty: 17 Professors, 9 Associate Professors, and 16 Assistant Professors.
COMPARE	The best typist in the pool can type 65 words per minute on the typewriter and 80 words per minute on the word processor.
RESTATE	In our office, the new copier produces copies 2.5 times faster than the old copier. The old copier produced 50 pages every minute.
VARY	An editor of a prestigious journal noticed that, for a particularly wordy author, there were five reference citations for every page of text. There were 35 text pages in the manuscript.

individuals learn new concepts that have obvious ties to much of their previous knowledge.

The Nature of Instruction

To model successfully the acquisition of knowledge from instruction, one must examine the nature of that instruction and the type of information contained in it. Generally, there are two ways to present new concepts to students. The instructor can introduce the name of the concept and give a prototypic example. The example contains specific details and is couched in a setting that should be well-understood by students. An alternative approach is for the instructor to provide the name of the concept and give a general description of its most important features. This information is abstract and contains basic characteristics that should apply to all possible instances of the concept. In practice, instructors typically do both. They introduce a new concept by name, give a representative case in which the

concept clearly occurs, and then make a broad statement about the concept, which is intended to help the learner generalize the concept from the given example to other potential instances.

Some interesting research has been carried out to determine whether students learn differentially under different instructional conditions. Usual studies of instructional content tend to contrast one form of information with another, such that each student sees only one type. An example of this type of research is found in Sweller's (1988) comparison of problem-solving performance following rule-based or example-based instruction.

The issue I address is different: Given access to typical instruction in which both specific information (i.e., examples) and abstract information (i.e. definitions) are available, which will a student remember? Do students commit equal amounts of specific and abstract knowledge to memory? Is one type necessarily encoded first, to be followed by the other? Are there large individual differences? If so, are these differences related to performance? The following experiment provides some initial answers to these questions.

THE EXPERIMENT

Subjects

Subjects were 27 college students with relatively weak problem-solving skills. They were recruited from introductory psychology classes. On a pretest of ten multi-step arithmetic word problems, they averaged six correct answers.

Procedure

Each student worked independently on a Xerox 1186 Artificial Intelligence Workstation. All instruction and exercises were displayed on the monitor, and the student responded using a three-button optical mouse. Each student participated in five sessions, with each session comprised of computer instruction, computer exercises, and a brief interview. Students spent approximately 45-50 minutes working with the computer in each session and talked with the experimenter for about 5-10 minutes in the interviews. As stated previously, only the first session--the introduction to the five situations--is of interest here.

Data Collection

Data were collected from two sources: student answers to the first exercise presented by the computer and student responses to the interview questions. Each is described below.

Identification task. The first source of data was the computer exercise that followed the initial instructional session. The items in this task resembled those of Table 1. They were selected randomly for each student from a pool of 100 items, composed of 20 of each type. During the exercise, one item at a time was displayed, and the student responded to it by selecting the name of one situation from a menu containing all five names: *Change, Group, Compare, Restate, Vary*. The student received immediate feedback about the accuracy of the answer, and if the student responded incorrectly, the correct situation was identified.

The order of item presentation was uniquely determined for each student. Items of each situation type remained eligible for presentation until one of two criteria was obtained: Either the student had given correct responses for 2 instances or the student had responded incorrectly to 4 of them. Thus, a student responded to at least 2 items of each type and to no more than 4 of them. The minimum number of items displayed in the exercise for any student was 10, which occurred only if the student answered each of them correctly. The maximum number that could be presented was 20 items, which could happen only if a student erred in identifying the first two items of all five types. The number of items presented ranged from 10 to 18.

Interview Responses. The second source of data was information given by the students in the interviews. The interview followed immediately after the identification task described above. During the interview, each student was asked to describe the situations as fully as he or she could. The student was asked first to recall the names of the situations and then to describe each one that he or she had named. After each of the student's comments, the experimenter prompted the student to provide additional details if possible. All interviews were audiotaped and transcribed.

It is the interview data that reveal which pieces of instruction were encoded and subsequently retrieved by each student. Certainly, not all of the new knowledge acquired by an individual will be revealed in an interview. It is expected that students have more knowledge than they can access (as pointed out by Nisbett & Wilson, 1977). Nevertheless, the interview data are indicative of how the individual has organized his or her knowledge of the newly acquired concepts, and they suggest which pieces

11

of knowledge are most salient for the individual. Following well-known studies such as Collins and Loftus (1975) or Reder and Anderson (1980), we may assume that individuals will tend to retrieve the most closely associated features and those with highest salience for the individual.

Knowledge Networks and Cognitive Maps

Data from the student interviews were used to construct knowledge networks, one for each student. Each network consists of a set of *nodes*, representing the distinct pieces of information given by the student, and *links* connecting the nodes, representing associations between the pieces of information.

The interviews were coded in the following way. First, irrelevant comments were eliminated. These were things such as "Um, let me think" or "I'm trying to remember ...." Next, distinct components or elements of description were identified. These were usually phrases but could also be single words. These became the nodes of the knowledge networks. Two nodes were connected in a network if the student linked their associated pieces of information in his or her interview response. Two research assistants and the author coded each interview with complete agreement.

In addition to the knowledge network for each student, an "ideal" network was constructed from the instructional text. As with the students' networks, nodes were created to represent each distinct piece of information. Two separate pieces of information appearing contiguously in the text were represented by two nodes with a link between them. Needless to say, this network was substantially larger than any student network. It represents all that a student could possibly encode from the instruction, and thus it serves as a template against which to measure the amount and type of information encoded by each student. The "ideal" network for all of the situational information is presented in Figure 1.

Two things should be noted about the network presented in Figure 1. First, distances between nodes and spatial orientation of the nodes have no meaning. Only the presence or absence of nodes and links is of importance. Second, in this figure, all nodes appear equally important, and the same is true for the links. Strength and activation are not shown. However, in theory each node has a measure of strength that is a function of how many times it appears in the instruction, and each link has a similar measure of activation, depending upon how frequently the two nodes are linked.

12

Figure 1: THE "IDEAL" NETWORK

Figure 1 represents the ideal case in which all information is included in the network. Normally students do not retain all of the details, and the networks one constructs for them appear incomplete when compared with the ideal situation. Thus, we expect the student networks to be considerably sparser than that shown in Figure 1.

Several types of information may be gleaned from a student's knowledge network. First, of course, the network is an indication of how much the student remembered. The number of nodes in a network provides an estimate of this information. Second, the network shows which pieces of information are related for an individual. A measure of association can be made by counting the number of links and using that number to estimate the degree of connectivity of the entire network. Node count and degree of connectivity are standard network measures. I have discussed elsewhere how they may be used to estimate a student's knowledge of a subject area (Marshall, 1990).

13

In this chapter I examine two additional types of information: (a) *specificity*, which is the students' tendency to recall specific or abstract features to describe the situations and (b) *confusions*, which show the extent to which students confused different aspects of the five situations. One examines nodes to estimate the former and links to estimate the latter.

Specificity. Each node in the "ideal" network reflects one of two types of detail: specific or abstract. Specific knowledge refers to elements of information having to do with the examples presented in instruction, and it reflects the particular details of the example. Abstract knowledge refers to the general features or definition of the situation. The instruction contains approximately an equal amount of both types, as can be seen in Figure 1. The abstract nodes are represented by filled circles, and the specific ones are indicated by hollow circles.[4]

Each distinct piece of information (i.e., each node) recalled by a student was categorized as being *specific* or *abstract*. A response was considered to be *specific* knowledge if it pertained to a specific example. Typically, students giving this sort of response referred to details from the initial example used in the computer instruction. An illustration is given in the specific response of Table 2. The italicized phrases are examples of specific detail. In contrast, a response was considered to be abstract knowledge if it reflected a general definition or characterization. Table 2 also contains an illustration of an *abstract* response, and the italicized phrases indicate the abstract detail. The final example of a student response in Table 2 illustrates the case in which neither abstract nor specific detail is recalled.

Three measures of specificity were developed: the number of specific responses, the number of abstract responses, and the ratio of abstract to specific responses. These measures were used in the statistical analyses described below.

Confusions. In the networks representing situational knowledge, two types of links are possible, intra-situational and inter-situational links. Intra-situational links are judged always to be valuable. That is, if two nodes are both associated with one situation and they are connected to each other, then the retrieval of one of the nodes ought to facilitate the retrieval

---

[4] It should be noted that the instruction was not developed under the constraint that equal abstract and specific details be contained in it. The guiding principle was to explain each situation as completely as possible, using specific and/or abstract elements as needed.

Table 2
Examples of Student Responses

ABSTRACT	Q:	What do you remember about Group?
	A:	Group is when you have different items, *different groups of items*, that can be *categorized into one general group*.
SPECIFIC:	Q:	What about Group?
	A:	That was when you bought *7 shirts and 4 pairs of shorts* and they grouped it into clothing. So you had *11 separate things of clothing*.
NONE:	Q:	Tell me about Change.
	A:	I pressed that review button so many times and I can't remember anything right now. Um, change was, um my mind is blank right now. I did okay on the computer. I've forgotten just about everything. I'm trying to think of an example. I know they change something and make something else.

of the other. This is the principle of spreading activation. In general, the more knowledge the individual has about a concept and the greater the number of associations connecting that knowledge, the better the individual understands it. Figure 2 shows how the "ideal" network of Figure 1 can be represented as a two-layer map. The nodes at the upper level are the five situations, and those at the lower level are the knowledge nodes developed during instruction. Connections among the nodes at the lower layer represent intra-situational links. Generally, a larger number of connections at this level indicates greater understanding on the part of the individual. It is these connections that are shown as well in the network of Figure 1.

In contrast, inter-situational links, i.e., links between different situations, may or may not be of value to the individual's learning, because they are a potential source of confusion. Such links will not always reflect confusions; situations could in principle share one or more features. In the present case, however, the instruction was carefully designed to eliminate common features among situations. This is reflected in Figure 2 by the connection from each node at the lower level to a single node at the upper

Figure 2: THE "IDEAL" MAP

level. Given the design of instruction, there should be no inter-situational links. That is, no node at the lower level should connect to more than a single upper level node. Such linkages would be confusion links and reflect a misunderstanding about the two situations so linked.

An example of differences in students' inter-situational and intra-situational links is given in Figure 3. Two student maps are presented in this figure. Both students encoded a relatively large amount of information from the instruction, compared with other students in the experiment, but it is clear from the figure that they recalled different elements of information. Student S7 remembered distinct pieces of information about each situation and showed no confusions. S22, on the other hand, expressed a number of confusions, which are represented in Figure 3 by the dashed links between the two layers of nodes. These cognitive maps are characteristics of incomplete mastery. The situational knowledge of every student can be described by such a map. Obviously, the deficits of a student are highly individual. These individual differences will be discussed further in a later section of this chapter.

In summary, the student network and its corresponding map provide information about the number of details the student remembered about a situation, the amount of connectivity, the type of knowledge (i.e., abstract or specific), and the number of confusions in the student's response. The networks and the measures described here were the bases for the statistical

Figure 3: TWO STUDENT MAPS

analyses presented below and also served as input to the simulation model, which is described in the section following the statistical analyses.

STATISTICAL ANALYSIS

Three questions are addressed by the statistical evaluation. The first is whether students remember different amounts of detail from instruction, the second is whether one can characterize the type of information encoded by a student, and the third is whether these differences are related to the students' success on the identification task. Evaluation of the student networks shows that some students were more likely to encode mostly specific details, some were more likely to encode mostly abstract information, some encoded both in about equal proportions, and some encoded almost nothing. The statistical analysis evaluates whether these tendencies are related to performance on the identification task and whether the relationship can be generalized to the entire group of students.

It is evident from the interview data that students varied greatly in the amount of information they were able to recall about the five situations. The number of different details retrieved by students extended from a low of 3 to a high of 20. The mean number of details was 13.5, with a standard deviation of 4.02.

The number of abstract and specific details recalled also varied, and the ratio of abstract to specific detail ranged from 14:3 to 6:14. Thus, the answers to both the first and the second questions are affirmative: There were clear differences in the total amount of information recalled as well as differences in the amount of abstract and specific information.

Two analyses provide insight into the importance of this difference. First, on the basis of their interview responses, students could be divided into three groups: *Abstract, Specific,* and *Both.* Students classified as *Abstract* gave predominantly definitional responses in the interview. Those classified as *Specific* used mostly example information from the computer instruction to describe the situations. Those classified as *Both* responded with approximately equal numbers of abstract and specific detail. For membership in either the *Abstract* or *Specific* group, students had to have given at least 9 different pieces of information during the interview with at least twice as many instances of one type of information as the other. Approximately equal numbers of students could be classified as *Abstract* or *Specific,* with 6 in the former and 7 in the latter. An additional 11 students were categorized as *Both.* These students gave at least 9 responses with approximately equal numbers of abstract and specific details.

Figure 4 shows the relative performance on the identification task of the three groups described above. A one-way analysis of variance, with a dependent measure of correct responses to the identification task,[5] indicates that the groups differed significantly in their ability to recognize the situations, $F(2, 21) = 4.53$, $p < .025$.[6] As can be seen in Figure 4,

---

[5] It will be recalled that students viewed differing numbers of items on this exercise. For purposes of comparison in this analysis, only the first two exemplars of each type of situation were scored. Thus, each student received a score from 0-10.

[6] Complete data were not recorded for two students. One loss was the result of computer failure and the second was the result of a malfunction in the recording of the interview. These two students were excluded from the analyses reported here. Two other students having only 6 and 3 interview responses respectively were also excluded from this analysis.

Figure 4: GROUP PERFORMANCE

students who responded primarily with abstract characterizations of the five concepts were most successful, followed by those who used both types of information. The group relying on examples only were less successful than those using abstract only or abstract knowledge in conjunction with specific details. The performance of the abstract group was significantly higher than the performance of the example group, $t(21)= 3.005, p < .01$.

The above analysis shows that differences in student performance can be explained in terms of whether a student remembered abstract or specific information. One also expects that the absolute number of details that a student remembers--regardless of whether they are definition or example-- would be a good predictor of performance. Surprisingly, this is not the case. The Pearson product moment correlation between the number of correct responses on the performance test and the total number of nodes encoded from the student's interview is .074, accounting for less than 1% of the variance.

A second and more informative way of analyzing the data is a multiple regression analysis based on the type and amount of information, the inter-situational confusions, and the interaction between the two. In this analysis, the predictors are (1) X_1, the ratio of abstract to specific detail, (2) X_2, the number of confusions mentioned explicitly by the student, and (3) X_3, a product variable of the first two predictors. The dependent measure, again,

is the 10-item identification task. The resulting prediction equation was: $Y' = 6.667 + .602X_1 + .545X_2 - .617X_3$, with all coefficients reaching the conventional .05 level of significance. The model accounted for 43% of the variance and was statistically significant, $R^2 = 0.43$; $F(3, 21) = 5.38$, p < .01.

In general, students with higher abstract to specific ratios performed better on the identification task and made fewer confusion errors. Students with low ratios (i.e., those with more specific answers) named relatively few confusions but also responded with fewer correct answers. Students with approximately the same number of specific and abstract responses had the greatest number of stated confusions.

Thus, the statistical analyses suggest several group characteristics with respect to learning new concepts. That is, there are tendencies of response that apply over many individuals, not just a single one. These analyses are based on summaries of the cognitive maps and aggregate responses to the identification task. A more detailed investigation of individuals' responses provide additional information about the nature of learning in this study.

THE COGNITIVE MODEL

A more exacting analysis of the relationship between each student's cognitive map and his or her responses to the identification task was carried out by simulating the responses using a simple feed-lateral connectionist model. The model simulates for each student his or her response to each item of the identification task that the student actually attempted to identify.

The general model is given in Figure 5. It has three types of units: inputs, student nodes, and outputs. Inputs to the model are coded representations of the problems, and outputs are the names of the situations. As in most connectionist models, activation spreads from the input units at the lowest level to those of the intermediate level(s) through their connections. At the middle level, activation spreads laterally from the nodes directly activated by the lower level units to other nodes at the same level with which they are linked (this is represented by the two middle layers in Figure 5). Finally, the total activation coming into each unit at the output level is evaluated, and the output unit with the highest activation is the model response. Unlike many connectionist models, the units at the middle layer, their connections with other nodes at this level, and their linkages to the upper level are determined explicitly from empirical data.

Figure 5: THE FULL MODEL

21

Table 3
Item Characteristics Used to Encode Story Situations

General Characteristics:
Set modification
Permanent alteration
Class inclusion (explicit or implicit)
Relation between two objects
Relation between an object and a property of that object
Fixed relation (implied)
Relative size
Size differential
Percentage
Causality
Multiple agents
Multiple objects
Unit measurement
Two identical relations

Key Phrases:
Each/every/per
As many as
Have left
Altogether/A total of
More/less
Cost
Same
If ...Then
Money

Time Features:
Specific time elements (minutes, days, weeks)
Before/after

The bottom layer of units. The inputs consist of information about the items that comprise the identification task. There are 27 possible characteristics that can be present in any item. The set of characteristics is given in Table 3. Each item is coded according to these characteristics as a 27-element vector containing 0's and 1's, with 1 indicating the presence of a

characteristic and 0 its absence. Not all characteristics will be present in any single item; usually a simple situation requires only a few of them. The mean number of characteristics for the 100 items used in the identification task was 4.33. All 100 items were encoded by three raters with complete agreement.

The middle layers of units. For each student, the middle layers of the model contains a set of nodes and the connections between them. The two layers have identical sets. The nodes and links were identified from the student interviews, as described previously, and they formed the basis of the statistical analyses of the preceding section. Three trained individuals read the transcript of each individual's interview and determined which nodes were present and whether they were linked. As in the characteristics coding above, the three coders were in complete agreement.

The top layer of units. The outputs for the model are the five situation names: Change, Group, Compare, Restate, and Vary. Only one output is produced for a given input vector. The five possible outputs compete, and the one with the highest accumulated activation wins.

Connections between the bottom and middle layers. Each input element may connect directly to one or more of the nodes contained in the student's network (represented by the middle layers of nodes). Two layers are needed in this model to illustrate the feed-lateral aspect. The lower of the node layers connects to the input units. The second layer illustrates how the nodes connect with each other. Each node from the lower node set connects to itself and to any other nodes to which it is linked, as determined from Figure 2. Thus, activation spreads from the input units to the lower node layer. Each node transfers its own activation to the next layer and also spreads additional activation to any other nodes to which it is connected. This particular two-layer representation of a feedlateral network preserves the usual constraint that activation spreads upward through the model.

Some of the input elements (i.e., those units represented at the very bottom of Figure 5) may activate many nodes in the network, some may activate only a few, and some may fail to make a connection (if the student lacks critical nodes). The allowable linkages between the input and middle layers of units were determined by mapping the input characteristics to the "ideal" map of the entire instruction. Recall that the input characteristics are general features. Most of them activate multiple nodes, and these nodes are frequently associated with different situations. Thus, it is rare that one input characteristic points to a single situation. The full pattern of possible activation is shown in Figure 5. Note that this figure illustrates all

characteristics as they link to all nodes and is thus a theoretical pattern. The model would never be presented with a problem containing all possible features, nor did any student have all possible nodes at the middle layers.

Once the student network receives the input, activation spreads from the nodes directly targeted by the input elements through any links they have to other nodes at this level. All of the activated nodes then transmit their total activation to the units at the upper level. The amount of activation for each situation is determined from the accumulation of activated links leading to it. The five situations compete with each other for the highest level of activation, and the one with the highest value becomes the output. Thus, the model of Figure 5 represents the input of an item, the activation of the student's semantic network, the competition among situations, and the final output as a result of total activation throughout the model.

The model depends upon the set of nodes for each student, the pattern of linkages among them, the overall association of subsets of nodes with the situation labels, and the input characteristics of the items. All except the latter are derived from the student cognitive maps described earlier.

Model Verification

As a test of the model's adequacy, a simulation was carried out in which the ideal network of Figure 2 was used as the student model. The 100 items available in the identification task were presented to the model, and its responses were compared with the correct answers. The model performed with 100% accuracy, successfully identifying the situations for all items.

Simulation Results

A simulation of each student's performance on the identification task was carried out. For each student, the response to the first item encountered in the exercise was simulated first, using that item's vector of characteristics and the student's network information. The second item followed, and then all subsequent items until the exercise terminated. Thus, the simulation covered all items presented to the student in the order in which the student saw them.

As described above, the number of items answered by students varied from 10 to 18, yielding a total of 360 item responses. A comparison of the results of the simulation of these 360 responses with the actual student responses to them is given in Table 4.

Table 4
Simulation Results

Outcome	Frequency Observed Outcomes	Frequency Adjusted Outcomes*
CSM	192	192
C_SM	64	64
C_S_M	19	13
CS_M	30	13
CM_S	55	51
Total	360	333

Key:

(1) CSM Both model and student answered correctly.
(2) C_SM Model and student made the same error.
(3) C_S_M Model and student made different errors.
(4) CS_M Student answered correctly; model erred.
(5) CM_S Model answered correctly; student erred.

(C = correct response; S = student response; M = model response)

*Impossible matches excluded

Table 4 presents the observed classification of the students' responses as well as an adjusted classification against which the model was compared. All 360 items comprise the observed classification. In the adjusted classification, some items have been omitted from consideration because the model was constrained by a lack of information from the student interview. This occurred under the following condition: If a student was unable to remember the name of a situation or anything that described it in the interview, the model for that student would have no nodes at the middle layer that could link to the situation name. Thus, the model would be constrained to ignore that situation and would never generate a response pointing to it. Consequently, if a student omitted entirely a situation in the interview, all items for which the student gave that situation as a response were likewise eliminated. There were 27 of these impossible matches. As shown in Table 4, 17 of these were items which the student answered correctly, and 10 were items on which the student erred. It should be noted

that these are not model failures but are interview failures.

Each application of the model to a vector of item characteristics, representing a single item, resulted in one of five outcomes, as shown in Figure 5. Outcomes CSM and C_SM are exact, successful simulations of the model. In both cases, the model generated a response that was identical to the one produced by the student. In the first, the response was correct, and in the second, it was an error. The outcome C_S_M is considered to be a partial success of the model. Both the student response and the model response were in error, but they were different errors. In these cases, the model accurately predicted that the student lacked critical knowledge and would err.

The remaining two outcomes, CS_M and CM_S, represent simulation failures. The most serious of these is CS_M, reflecting cases in which the student answered correctly but the model failed to do so. They are serious failures because they suggest that the model did not capture sufficiently the student's knowledge about the situations. It should be noted that more than half of the observed instances of CS_M were impossible matches, as described previously. That is, the student omitted any discussion of the situation in the interview, and the model was subsequently constrained to ignore it. As mentioned above, these instances are considered to be interview failures rather than model failures. Only the remaining 13 instances are true model failures, representing just 3.9% of all responses.

The final outcome category, CM_S, also represents model failure but is less critical than the failures of CS_M. In this category, the model made a correct response when the student did not.

Many of the CM_S simulation failures can be explained by considering the students' experience as they respond to the identification task. During the actual task, many students made errors on one or more situations and then apparently learned to classify these same situations correctly. This is evidenced by their patterns of responses, typically an incorrect response to a situation followed by two correct responses to the same situation, with no additional errors. What has happened in such cases is that the student's knowledge network presumably changed during the course of the task. The knowledge base that generated the early incorrect responses is not necessarily the same one that generated the later successful ones. And, it is only the latter that is reflected in the student's interview. In such instances, the model would correctly match the two correct responses, but it would also give the correct response to the first item that the student missed. The model does not learn. It simulates the state of the student at the end of the

exercise, as reflected in the interview. If the student learned during the course of the exercise, we have no way of knowing what node configuration corresponded to the earlier, incorrect responses. Under the most conservative criterion of learning--an error followed by two correct responses--25% of the mismatches can be accounted for by student learning. In each case the model gave the correct response to all three items. [7]

Another 25% of the mismatches occurred when both the model and the student selected different wrong situations as the response option. In these cases, the model correctly determined that the student would not give the correct response. The model's answers may differ from the student's for a number of reasons, including guessing. These were, after all, multiple choice exercises, in which students were asked to select the correct situation from the menu of five possible ones. Students probably guessed at some of the answers, but the model does not guess.

There are other possible explanations for the model failures. On the one hand, some students may have been prone to "slip" as they made their selections using the mouse, resulting in the unintentional selection of the option residing either above or below the desired one. It is not an uncommon phenomenon, as those who use a mouse frequently can attest. Accidental errors of this sort are undetectable. Similarly, students may have used a test-taking strategy, such as avoiding the selection of one response if they used it on the immediately preceding exercise. These errors are also undetectable: The model does not take test-taking strategies into account.

If we consider the "probable learning" mismatches (i.e., those that were followed by two correct matches on the same situation) and the "different error" mismatches (i.e., those in which the model and student both erred but selected different errors) as *understandable or explainable* discrepancies, the total number of mismatches between students and the model is reduced from 77 to 51, leaving only 13 *CS_M* and 38 *CM_S* as mismatches. Thus, the model satisfactorily accounts for 85% of all student responses.

A final evaluation of the model's performance comes from examining how well individual student performance was simulated by the model. The

---

[7] Several other instances exist in which the student made multiple errors on a situation and then responded correctly to one final instance of that situation. While it is very plausible that learning also occurred in these cases, one hesitates to draw a conclusion based only on one response. Thus, these errors remain unexplained.

Table 5
Simulation of Individual Performance:
A Comparison of Model Responses with Observed Student Responses

Student	No. of Items	No. of "Impossible" Matches	Percent Exact Matches	Percent "Explained" Matches	Total Percent Matches
1	13	3	100%	0	100%
2	15	0	80%	7%	87%
3	13	3	90%	10%	100%
4	13	3	80%	10%	90%
5	16	0	75%	6%	81%
6	11	0	100%	0	100%
7	14	0	86%	7%	93%
8	13	0	92%	0	92%
9	14	5	100%	0	100%
10	15	2	85%	7%	92%
11	15	3	67%	8%	73%
12	16	0	63%	31%	94%
13	14	0	71%	0	71%
14	13	0	69%	8%	77%
15	14	3	100%	0	100%
16	13	3	90%	10%	100%
17	15	0	73%	14%	87%
18	14	0	79%	7%	86%
19	14	0	79%	7%	86%
20	18	0	67%	7%	72%
21	13	0	69%	0	69%
22	16	0	69%	12%	81%
23	16	0	56%	13%	69%
24	16	2	57%	7%	64%
25	16	0	69%	12%	81%

results for each student simulation are given in Table 5. Two measures of success are given in the table. The first is the number of exact matches, excluding the "impossible" ones. The second is the overall percentage of satisfactory matches for each individual and is given in the extreme right-hand column of the table. This percentage is based on the number of satisfactory matches, including the "probable learning" and "different error" mismatches described above but eliminating from consideration the "impossible" matches. As can be seen in Table 5, the performance of 6 of the 25 students was fit exactly by the model with 100% agreement. The

model simulated the performance of an additional 12 students with accuracy between 80-99%. The model's success rate fell below 70% for only 3 students, to a low of 64%.

DISCUSSION

There are several important implications that result from this study. They are discussed below with respect to the three questions posed in the introduction: What do they learn, when do they learn it, and what can they retrieve?

What specific information does a student learn from initial instruction about a new topic?

One of the most striking findings was that students tended to encode and use specific details from the initial examples used in instruction. Almost all of the example nodes had to do with the five introductory examples, despite the fact that several other examples were given later in the instruction. (See, for example, the *Specific* response of Table 2.) This finding suggests that the very first example of a concept is highly important and should, therefore, be carefully developed. For many students, the initial examples provided the scaffolding for the semantic networks. Some of the details of those examples led to erroneous connections. As a case in point, the example for one of the situations was based on money, leading some students to expect (incorrectly) this situation to be present whenever money was in the problem. These faulty connections were very evident in their interview responses.

A general pattern of encoding was apparent from the students' responses. Several students described the situations only in terms of the examples. When prompted, they were unable to embellish their descriptions by using abstract characterizations. No instance of example information followed by abstract information was observed. In contrast, students having abstract knowledge always used it in preference to giving example details. That is, their initial responses were generalizations. When prompted for more information, they used example details to support their abstract descriptions. These findings suggest that students may first encode the example information and then build the abstract network around it. Once formed, the abstract portion of the network becomes stronger upon exposure to additional examples, whereas the example portion does not augment its activation or strength. If the abstract information is not encoded, the details

of the example—which received high strength initially—remain the most salient elements of the network.

How is the information that the student encoded in memory related to the student's performance?

The statistical analyses suggest that the degree to which a student is able to use his or her abstract information is positively related to the student's success on the identification task. Those able to express mainly abstract knowledge apparently had the best understanding of the five concepts and were most easily able to identify them. Those for whom the abstract characterizations were somewhat incomplete (e.g., those who were able to give abstract description for some concepts but needed example details to describe others) performed less well but still were more successful than those who predominantly relied on example details.

The primary implication of this finding is that instruction should be developed to facilitate the linkage of abstract knowledge to easily understood example knowledge. The examples were undoubtedly salient and easily encoded. For some students, the abstract characterizations were equally easy to encode, but this was not universally true.

Does the cognitive model reflect this relationship?

The connectionist model is a useful way to examine individual performance of students as they identified these concepts. The simulation of individual performance was extremely successful. The high level of agreement between model performance and student performance suggests that the model captures most of the salient and discriminating information actually used by the students. Most important, the model demonstrates the impact of missing nodes and erroneously linked pairs of nodes. In many cases, knowledge of which nodes were missing led to accurate predictions of subjects' erroneous responses. In others, incorrect linkages among nodes also led to accurate predictions of errors. The model and its simulation provides strong support for the use of cognitive networks to represent learning of concepts.

REFERENCES

Collins, A. M. & Loftus, E. F. (1975). A spreading activation theory of semantic processing. *Psychological Review, 82,* 407-428.

Marshall, S. P. (1990). Selecting good diagnostic items. In N. Frederiksen, R. Glaser, A. Lesgold, & M. Shafto (Eds.) *Diagnostic monitoring of skill and knowledge acquisition* (pp. 433-452). Hillsdale, NJ: Lawrence Erlbaum Associates.

Marshall, S. P., Barthuli, K. E., Brewer, M. A., & Rose, F. E. (1989). *STORY PROBLEM SOLVER: A schema-based system of instruction.* Technical Report 89-01 (ONR Contract N00014-85-K-0661). San Diego: San Diego State University, Center for Research in Mathematics and Science Education.

Marshall, S. P. (1991). *Understanding the situations of arithmetic story problems: A basis for schema knowledge.* Unpublished manuscript.

Nisbett, R. E. & Wilson, T. D. Telling more than we can know: Verbal reports on mental processes. *Psychological Review, 84,* 231-259.

Riley, M. S., Greeno, J. G., & Heller, J. (1983). Development of children's problem-solving ability in arithmetic. In H. Ginsburg (Ed.), *The development of mathematical thinking* (pp. 153-196). New York: Academic Press.

Tarmizi, R. A., & Sweller, J. (1988). Guidance during mathematical problem solving. *Journal of Educational Psychology, 80* pp. 424-436.

Problem-Solving Schemas: Hybrid Models of Cognition

Sandra P. Marshall
San Diego State University
San Diego, CA

John P. Marshall
Crystal Graphics, Inc.
Santa Clara, CA

An ongoing controversy in the cognitive science community centers on the nature of the models used to represent cognitive phenomena. The two primary contenders are production-system models (such as ACT* and SOAR) and connectionist models (such as those produced by McClelland and Rumelhart or Grossberg and his associates). Critics of both sides argue that the other cannot suffice to capture human behavior. Both appear to be right. Perhaps what is needed is a model that combines the best--and lessens the worst--features of both kinds of model. A hybrid model having these characteristics is the topic of this report.

The recognition that hybrid models are needed is not new. A number of prominent researchers (from both sides of the argument) have suggested that some union of the two representations is in order. For example, in *The Computer and the Mind,* Philip Johnson-Laird hypothesized that one way to get around some of the dilemmas posed by existing models of cognition was to "postulate different levels of representation: high-level explicit symbols and low-level distributed symbolic patterns" (p. 192). The opinion of a long-time connectionist is reflected in the title of a recent article: "Hybrid Computation in Cognitive Science: Neural Networks and Symbols" (J. A. Anderson, 1990). And, Marvin Minsky echoes the sentiment in his 1991 paper, "Logical Versus Analogical or Symbolic Versus Connectionist or Neat Versus Scruffy," where he states explicitly that we "need integrated systems that can exploit the advantages of both" (p. 37).

The need to combine the two representations derives from the fact that neither alone has been entirely satisfactory in modeling complex cognition. Symbolic production systems, as the oldest and most widely used of the two, have been very successful in describing some important aspects of rule-based problem solving. They have been widely used in artificial intelligence and have greatly influenced the development of intelligent tutoring systems (Winger, 1987).

At the same time, such systems are noted for their inflexibility on some relatively simple tasks, such as object recognition and classification. Bereiter (1991) provides a good discussion of some of the central problems with rule-based

cognitive models. As he (and others) point out, humans are not particularly good at working out extended logical sequences. They make mistakes. A production system does not make mistakes, and one difficulty in working with production models is to produce human-like errors from the model. Many models consistently have better performance than the humans they are intended to mimic.

A second difficulty lies in the way that production systems work, namely systematically, orderly, and efficiently. People don't seem to have those characteristics. We see this difficulty when we try to model complex problem solving, using protocols generated by experts or by novices. Very few individuals start at the beginning of a problem and proceed carefully through a top-down process to reach the solution. To model their performance, we all too often are forced to disregard some of the protocol material in our quest for sequential rule-based performance. Moreover, many individuals simply cannot articulate what they are doing or explain why one part of a problem triggers a particular response from them, which suggests that their activity is not entirely a neat and orderly process.

Nonetheless, there are clearly many instances in which individuals do engage in rule-based cognition, and production systems have to date provided our best means of modeling them. This is particularly apparent in well-specified domains from mathematics such as arithmetic, algebra, or probability, and in areas of physics such as electricity and magnetics. What is common in these domains is that there are highly specific rules that need to be acquired and applied by individuals in order to operate successfully in the domain. As a simple example, consider arithmetic operations. It would be the rare person who performed multiplication or long division without resorting to the use of a standard algorithm. Modeling the acquisition and use of such algorithms are precisely the areas in which production models excel.

On the other hand, connectionist models are weak in just these areas. Connectionist models excel in pattern recognition rather than in logical sequences of actions. Unlike production systems, connectionist models do not depend upon the firing of independent units such as rules. Rather, a collection of units (nodes) spread activation through their connections to other units. One does not trace the history of a cognitive process very easily in a connectionist model because of this feature. Subtle differences in the connection weights may yield large differences in model response. At any point in the process, it is the pattern of weights that matters, not the presence or absence or a particular unit.

A particular strength of connectionist models is the flexibility allowed for inputs. Because the models depend on the of weights over a great many units, the presence or absence of any single unit is usually unimportant. Any input typically is characterized by a great many units.

Given the unique and complementary nature of the two approaches--the strength of the symbolic system for modeling sequences of actions and the strength of the connectionist approach for modeling pattern recognition--it is reasonable to anticipate hybrid models that will capitalize on their individual strengths. Surprisingly, few true hybrid models have yet emerged, although one suspects that the number under development is somewhat greater. One of the best examples available now was developed by Walter Schneider to model controlled or automated processes (Schneider & Oliver, in press).

The remainder of this report describes a particular hybrid model, a model of schema instantiation in arithmetic problem solving. This model utilizes both production systems and connectionist networks to represent schema knowledge.

Overview

Types of Schema Knowledge

Schema knowledge for problem solving consists of four major components: constraint knowledge, feature knowledge, planning knowledge, and execution knowledge (Marshall 1990; in press a, b, c). There are key issues involved in each type of knowledge, and each one demands its own distinct representation in the full model of a schema.

Constraint knowledge has to do with recognizing patterns. The question of interest is: Does the stimulus problem contain a pattern of elements sufficient to activate an existing schema? This pattern recognition is accomplished by a connectionist component of the model.

Feature knowledge, on the other hand, has to do with deciding whether the necessary elements are provided in the problem, *given that the pattern has already been recognized as characteristic of a schema,* so that the schema can be instantiated. This is a question best answered by a production system. There is also a connectionist part to this knowledge. Several potential patterns within one schema may exist in a problem, and the most reasonable or most likely one for solution needs to be recognized. This is a special case for competitive performance by all pattern candidates, to determine which pattern most strongly reflects the identified schema.

Planning knowledge is for the most part sequential and consists of setting goals and selecting operations for obtaining them. Again, a production system is appropriate. Planning knowledge guides the entire problem-solving process, and it calls on feature knowledge and constraint knowledge when it needs more detail or more elaboration about the problem.

Finally, execution knowledge involves the step-by-step execution of already-learned algorithms, which again calls for a production system. Execution knowledge comes into play only when the plans call for it.

These types of schema knowledge have been the focus of a number of experimental studies as well as the target of our modeling efforts. Each experiment typically spanned several weeks and required subjects to participate in a number of different tasks at various times. All of the experiments involved the *Story Problem Solver (SPS)*, a computer-based instructional program about arithmetic story problems, and/or the *Problem Solving Environment (PSE)*, a graphical system in which students could practice what they learned under *SPS*. These systems are described elsewhere (Marshall, Barthuli, Brewer, & Rose, 1989; Marshall, 1991).

Both *SPS* and *PSE* were designed around schema theory. In particular, they were developed so that each of the four components of knowledge described above could be isolated and evaluated as students acquired their schema knowledge. The results of the experiments using these systems are given in several other papers (Marshall, 1991; Marshall, in press a, c). The importance of the experiments for the present report is that they provide empirical evidence against which our computer models can be evaluated.

The Performance Model of Constraint Knowledge

We focused our attention initially on models of constraint knowledge. We did so for two reasons: First, problem solving typically involves two general aspects: recognition of the important parts of the problem and appropriate application of techniques to these components to obtain a solution. The recognition aspect demands constraint knowledge, suggesting that constraint knowledge is an important initial point of access to schema knowledge. Second, we were interested in how individuals understand and retain new information about a concept or set of concepts. This, too, falls under constraint knowledge.

Constraint knowledge can be modeled very well using relatively common connectionist models. Two models were created: a model that can mimic the performance of subjects and a model which learns when given appropriate feedback. Both of these were developed and evaluated as a first step in building the complete hybrid model.

The model that simulates subject performance in identifying the situations given in simple story problems is described in full in Marshall (in press)[1]. It is a

---

[1] The article referenced in this citation is reproduced in the chapter immediately preceding the present one in this final report. Details of the model are presented there.

three-layer feed-forward (and feed-lateral) model consisting of input units, a middle layer of interconnected units, and a set of output units.[2] The model takes as its input a binary vector having elements of 0 or 1, which represent 25 problem characteristics. The model produces as its output the identification of one of five situations that may occur in the story problem.

The input units are connected to a middle layer of units representing a subject's knowledge about the situations. This layer of knowledge nodes corresponds to the typical hidden unit layer found in many connection models, but it is not hidden in this instance. The nodes here derive from student interviews. Each student's interview about the situations was coded and transformed into a set of nodes and links among nodes. The model used this information to derive its output response. Again, full details of the simulation are given in Marshall (in press).

The performance simulated by the model is student response to a computer-based identification task. The model performed very well, simulating the performance of a number of students exactly and accounting for a large majority of responses for the rest. Both correct responses and specific errors were modeled.

The Learning Model[3]

The performance model provided the initial framework for the subsequent learning model. The input units for this model are essentially the same as for the performance model, as are the output units. A layer of hidden units replaces the knowledge nodes that derived from the student interviews in the performance model. Thus, these hidden units are hypothesized to exist but are unknown. Moreover, the nature of their connections to the inputs and outputs are unknown. The question of interest in this model is whether a connectionist model of this form can learn to make the appropriate classification of the five situations.

The optimum number of hidden units for this case is undetermined. Theoretically, it depends upon the optimum number of knowledge nodes that a student should acquire, and this number is not known. From the instructional

---

[2] The model can also be conceived as a four-layer feed-forward model in which the second and third layers contain the same units. This eliminates the problem of having activation spreading laterally among units at any level. To achieve independence at all levels, we insert a third layer of units that duplicates the units at the second layer. Connections exist from all original units at the second layer to their counterparts at the third layer and also to any other units to which they may be conceptually related.

[3] In developing this model, we have drawn substantially from the models described in chapter 8 of Rumelhart & McClelland (1986) by Rumelhart, Hinton, & Williams as well as the extension of them in chapter 5 of McClelland & Rumelhart (1989).

experiment reported in Marshall (in press), the maximum number would be 33, the number of possible knowledge nodes. However, no student ever acquired all possible nodes, and it is not clear that having all of them would produce maximum performance. Several students made more than 90 percent successful responses with many fewer nodes. In *SPS*-based experiments, we observed that students typically acquired an average of 14 nodes. The range was 6 to 17. In general, having more nodes did not necessarily mean that students performed more successfully on the task.

Using the instructional experiment as a guideline, we include 14 hidden units in the learning model. As an initial simplification, the model is constrained to have 3 layers, eliminating the feed-lateral feature of the performance model. The three layers are the input layer, containing information about the problem to be classified; the hidden layer, containing units that correspond roughly to the knowledge nodes of the performance model; and the output layer, containing the names of the five possible situations.

The model requires specification of a learning rate, η . This rate defines how strongly the model reacts to incorrect answers with each trial. The learning rate must be chosen carefully so that the system will converge to the correct solution in a reasonable amount of time. A learning rate which is too large may cause the system to converge to an incorrect solution while a learning rate which is too small may prevent convergence in a practical amount of time (and perhaps at all). In general, for the network to stabilize (i.e., for learning to occur), the larger the number of hidden units, the lower the learning rate. We found that learning rates between .05 and .10 were most satisfactory.

For this model we also include a momentum factor, μ . This factor allows the system to carry over learning from previous problems when new problems are presented. As Rumelhart and McClelland (1989) point out, without the inclusion of a momentum factor, the system may converge to a "local" solution and stabilize there, even though there is a better "global" solution (p. 132). A suitably large μ prevents the model from getting stuck in such local solutions. In addition, a momentum term tends to speed up the model, because it allows the specification of a higher learning rate.

Testing the model corresponds to running it over enough trials for it to reach some pre-determined criterion. Each trial proceeds through the steps listed below:

   o presentation of a randomly selected input vector;
   o forward propagation of activation from input to hidden units
      and from hidden to output units;

o calculation of the errors associated with each output unit;
o backward propagation of errors from output to hidden units;
o modification of the weights of the connections between all
layers of units based on the errors.
Each of the main components of the model are described briefly below.

The Task. The task for which this model was developed is to learn the appropriate classification of a set of 100 story problems according to the situations depicted in them. Each problem is represented by a set of characteristics which the model uses in making its classification. Five output responses are possible.

Inputs. The inputs to the model are the set of 100 binary vectors nearly identical to the ones described for the performance model (see above and Marshall, in press). Each vector represents one arithmetic story problem. The problem is coded according to the presence or absence of several general characteristics.

The difference between these input vectors and those of the performance model is the inclusion here of coded information about the form of the question stated in the problem. In the performance model and in the empirical studies simulated by it, the items were situational descriptions and contained no questions. Both the learning model and the hybrid model described below require problems rather than situation if we are to model the full problem-solving process.

In general, there are two options for item presentation: either the entire set is presented again and again in some fixed order, making an orderly cycle through the entire stimulus list and insuring that each item is presented an equal number of times; or each presentation is randomly determined at the time of the trial, so that every item in the set has an equal chance of being selected on every trial. We have implemented the latter, primarily because we wished to avoid any possible order effects and also because items were always randomly generated for students in our empirical learning experiments.

Outputs. The model outputs correspond to the identifications of situations given in the story problems. For each problem presented, the model can make one of five possible responses, one for each situation.

Input Units. In a single trial the layer of input units is comprised of one input vector. Each element of the input vector takes a value of 1 (if the characteristic it represents is present in the selected item) or 0 (if it is absent). The input vector to the learning model contains the original 25 elements used in the performance model plus the additional 2 elements to code the question, resulting in a 27-element vector.

Hidden Units. The middle layer of the model contains hidden units. Each of these is connected to every input unit in the layer below it and each in turn

contributes to the activation of all output units above it. As mentioned previously, there are 14 hidden units.

Output Units. Each situation is represented by one output unit. On each trial, following activation from the hidden unit layer directly below, each output unit has some level of activation. The one with the highest level is the model response for that trial.

Bias. Each hidden and output unit has a bias associated with it. The bias is added to the incident activation upon the units and functions like a threshold for the unit (cf Rumelhart). If insufficient activation is received at the unit to overcome the effect of the bias, then the output of the unit will be insignificant.

Input-to-Hidden Weights. As in most connectionist models, each of the input units connects to each of the units in the layer immediately above it, i.e., the hidden unit layer, and each connection has its own unique weight. When an input vector is presented to the model, activation spreads from the input units to the hidden units. The amount of activation spread is determined in part by the strength (i.e., weight) of the connection.

Hidden-to-Output Weights. Each hidden unit is connected to every output unit, and each has a strength or weight. The values of these weights are also randomly generated for the initial trial, using the same constraints as for the input-to-hidden weights.

Model Parameters and Initialization Values

The model requires that two parameters be set: the learning rate η and the momentum μ. For most of our tests, we have used $\eta = .07$ and $\mu = .9$.

Additionally, the model requires that each unit i have a bias term β_i and that each connection between a pair of units i and j have a starting weight ω_{ij}. The bias terms and the weights are generated initially from a uniform distribution ranging from -.005 to +.005.

Finally, the learning criterion must be set. This requires choosing a tolerance value that indicates how many--if any--errors will be allowed and specifying how large the output value must be in order to be considered correct. We use a 90 percent tolerance standard; that is, the model must correctly identify at least 90 of the 100 test items. To be considered a correct response, the appropriate output unit must have a value that is at least .25 larger than the next largest output unit.

Under the parameter selections and initialization values described here, the model converges at approximately 7,000 trials.

Figure 1: The Connectionist Network

Technical Details of the Learning Model[4]

The model consists of the three-layer network shown in Figure 1, with each layer comprised of a set of units. Typically, one thinks of the input units as being at the bottom level of the model and the output units at the top, as in Figure 1. The hidden units make up the middle layer. Each layer is fully connected to the layers immediately above and/or below it, as shown in Figure 1.

We define the following elements of the model:

η	the learning rate;
μ	the momentum factor;
α_i	the activation that accumulates in each hidden or output unit i;
λ_i	the activation that spreads from unit i to units above it;
ω_{ij}	the weight associated with the connection between units i and j;
β_i	the bias associated with unit i ;
τ_i	the target level of output activation; externally set as 1 if the output unit i represents the correct situation or 0 if it does not;
ε_i	the error associated with unit i .

The model learns by processing an input vector and forward propagating activation from the lowest level to the highest, by calculating the error at this highest level and then backward propagating the error down through all levels, and by adjusting all weights connecting pairs of activated units accordingly.

The activation spreading out from a unit is defined as:

$$\lambda_i = \begin{cases} 1 \ or \ 0 & \text{if } i \text{ is an input unit} \\ \dfrac{1}{1+e^{-\alpha_i}} + \beta_i & \text{if } i \text{ is a hidden or output unit} \end{cases}$$

---

[4] All programs for the models in this report were written in C++ and run on a PC-80486 workstation.

where β_i is the bias associated with unit i and α_i is the activation that has accumulated into the unit from the layer of units below. The accumulated activation is determined by:

$$\alpha_i = \sum_{j=1}^{J} \omega_{ij} \lambda_j.$$

If i is a hidden unit, j refers to input units, and the summation occurs over all weights between a hidden unit and the input units at the level below and the λ_j's of the input units. If i is an output unit, j refers to hidden units. Each of the units j at the level below unit i will have an associated λ_j which influences unit i. Note that α_i is defined only for hidden and output units.

The spread of activation begins with the input units. Those with values of 1 activate their associated hidden units which in turn pass some of the activation to the output units (by means of their λ_i's). When the forward propagation of activation is completed, for every output unit i the difference between τ_i and λ_i is used to compute the error ε_i :

$$\varepsilon_i = (\tau_i - \lambda_i)\frac{\partial \lambda_i}{\partial \alpha_i}$$

where the derivative can be expressed as

$$\frac{\partial \lambda_i}{\partial \alpha_i} = \lambda_i(1-\lambda_i)$$

thus yielding a final form for the unit's error of

$$\varepsilon_i = (\tau_i - \lambda_i)\lambda_i(1-\lambda_i)$$

with all terms as defined above. The error signal is then passed to the hidden units by:

$$\varepsilon_i = \sum_{j=1}^{J} \varepsilon_j \omega_{ij} \frac{\partial \lambda_i}{\partial \alpha_i} = \sum_{j=1}^{J} \varepsilon_j \omega_{ij} \lambda_i(1-\lambda_i)$$

for each hidden unit i. The summation occurs over all output units j. Input units do not accumulate error.

After the error signal has propagated backwards through the network, the weights are adjusted by:

$$\omega_{ij}(t) = \omega_{ij}(t-1) + \Delta\omega_{ij}(t)$$

where t indicates the current trial and $(t-1)$ is the previous trial. The amount of change is determined by:

$$\Delta\omega_{ij}(t) = \eta\varepsilon_i\lambda_j + \mu(\Delta\omega_{ij}(t-1))$$

Note that w_{ij} emanating from input units with initial values of 0 will receive no adjustment. A similar adjustment is made for the bias terms, with

$$\beta_i(t) = \beta_i(t-1) + \Delta\beta_i(t)$$

and the amount of change is computed by

$$\Delta\beta_i(t) = \eta\varepsilon_i\lambda_i + \mu(\Delta\beta_i(t-1)).$$

After the weights and biases have been adjusted, the activation and error terms are reset to zero for the next trial. The only carryover from trial to trial is contained in the weights, the biases, and their delta values ($\omega_{ij}, \beta_i, \Delta\omega_{ij},$ and $\Delta\beta_i$).

The model runs with alternating learning and testing phases. The learning phase runs in blocks of 100 trials. At the conclusion of every block of trials, the model suspends the backwards propagation of error and runs a performance test over all input vectors to determine whether it has yet reached a specified criterion for successful learning. During a testing phase, the model maintains an unchanging set of weights, which is the set reached on the last trial of the previous learning phase.

The criterion for learning is the correct response to at least 90 of the 100 input items, with "correctness" established as the activation of the appropriate output unit i and with λ_i at least .25 larger than the next largest activation value for any output unit. In practice, the system typically converges with 94-97 items correct in the testing phase. Given the fact that λ_i ranges only from 0 to 1, a difference between two values of .25 is highly significant.

During the testing phase, each of the 100 input vectors is presented in a fixed order to the model and a response is generated following the spread of

activation as before. The response is scored as correct or incorrect, and the next vector is presented. If the model fails to reach the defined criterion (i.e., errs on more than 10 of the items), the learning phase resumes with another block of trials. When the model reaches the defined criterion, the weights are stored for later use, to be described below.

The Hybrid Model

The hybrid model developed to solve story problems has the form shown in Figure 2.[5] It has three main components: two production systems, represented in the figure by the decision boxes and arrows, and a connectionist network, represented in the figure as a set of nodes and links. All of these interact with each other, indicated in the figure by the arrows leading into and out of the rectangle in the middle of the figure. The connectionist model of Figure 2 is the identical model described under the learning model. It performs here in its testing mode; at this point the model is presumed to have learned the classifications, and no additional changes in weights occur. The weights used to compute a response for the hybrid model are those that were saved when the learning model reached its learning criterion.

The problem input to the full model now includes more than the vector of characteristics used in the performance and learning models. In addition to this vector—which remains the input to the connectionist part of the hybrid model—problem input consists of specific detail about the quantities found in the problem. This information is encoded by dividing the problem into several clauses. Each clause contains three types of information: owner, object, and time.

An example of clause coding for a specific problem is given in Table 1. *Owner* contains two fields: name and type. *Object* contains four fields: name, type, value, and action. The action will contain such information as is necessary to determine which arithmetic operation to use. For example, an action might be increase, decrease, more, or less. The final type of clause information is *Time*, which contains just one field that indicates a relative time of occurrence within the problem. A clause may contain multiple owners and multiple objects, and it may omit time. The clause information is provided as input to the smaller production system (indicated by the arrow in Figure 2).

---

[5] In this figure and others depicting hybrid models, we do not attempt to represent all units of the model. Rather, in the interest of having simple and easy-to-understand representations, we show only a few units at each level. Likewise, we do not show all rules that are part of the production system.

Figure 2: The Hybrid Model

When a problem is presented to the model, the connectionist network makes the appropriate recognition of the situation using the input vector as before, and it passes that information into a common area accessed by all parts of the model (represented by the rectangle). From here the information is available to the smaller production system (on the left side of the figure). This production system has as its goal the recognition of relevant elements of the problem once the situation is known. It passes its results back into the common area to be used by the connectionist network again if necessary or by the larger production system, which will produce a numerical solution.

The additional information derived from the clauses is used by the production system to determine which values of the problem are known, which are unknown, and their relationship to each other. Just as the recognition of the situation uses constraint knowledge about story problems, the selection of relevant pieces of the problem uses feature knowledge. To illustrate what we mean by feature knowledge, we describe briefly the change situation. A change situation is characterized by a permanent alteration over time in a measurable quantity of a single, specified thing. Thus, the model must confirm that only one thing is represented in the clauses. There are three aspects to a change situation: a starting amount, an amount by which it is to be changed, and an ending amount. The model must check that there are three available amounts, even if one of them is unknown. A change takes place over time, so the model looks for three distinct times to be represented in the change situation. The production system works through the clauses, confirming that similar elements are involved and placing the values from the problem on a list that can be used by the larger production system in the model.

Thus, for the hybrid model of Figure 2, the constraint knowledge is modeled by the connectionist network, and the feature knowledge is modeled by a production system. In the full hybrid model of schema knowledge, relevant feature and constraint knowledge are used to plan a solution. Thus, the input to the planning component of the hybrid model is the output from the feature production system coupled with the output from the connectionist model of constraint knowledge. Together, they provide sufficient information for the planning production system to set a series of goals and to call on the appropriate execution knowledge for achieving them. Table 1 illustrates some of the production rules.

Of the four types of knowledge that comprise a schema, we consider execution knowledge to be the least interesting, and we have made little attempt to model how individuals learn the basic arithmetic operations. We take as given that these are in place. Our argument here is that there already exist production systems designed to model the acquisition and use of the algorithms of addition, subtraction,

multiplication and division. The model here focuses instead on the selection of the appropriate values from the problem to use in carrying out necessary computations. Thus, the errors that can be modeled are those reflecting mistakes in selecting pieces of the problem or in selecting an operation to be carried out. Errors of computing are not possible (i.e., 3 x 4 = 7). The consequence of this assumption about algorithms is that we have not constructed a separate production system to make the computations, although it would be easy to do so.[6]

The model is instantiated with input to the connectionist model in the lower portion of Figure 2. The input consists of a single binary vector representing all information in the multi-step problem. Thus, pointers to more than one situation typically occur. The connectionist network identifies the most salient situation, and passes that information to the feature identifier (represented in Figure 2 as the smaller of the two production systems). Using the output from the connectionist network together with the clause information, this part of the model determines the best configuration of data to represent the selected situation. Several configurations may be possible. The production system selects a subset of clauses to represent each one. If there are multiple configurations, each one is then evaluated using the original connectionist model. For each configuration the production system creates a new input vector that contains only the information of the selected clauses. The output values associated with each input vector are compared, and the input vector leading to the highest value is selected as the immediate problem to be solved. The identified situation and its subset of clauses are then passed to the planning component of the schema. Thus, feature knowledge (i.e. the production system) and constraint knowledge (i.e., the connectionist network) interact to provide the necessary information that will be used to plan the solution.

A plan begins with the creation of a goal stack in which the top level goal is to produce a numerical solution. Additional goals are added to the stack and removed as they are achieved. A number of different goals are addressed by the production system. Some have to do with locating the unknown in the problem. Others center on carrying out the appropriate computations. Like feature knowledge and constraint knowledge, planning knowledge is schema-specific. The model uses its knowledge about the current schema to develop plans for solving the problem.

Table 1 illustrates a number of different goals and the steps the model takes to achieve them. The model attempts to solve the first subproblem it recognizes. If it is successful at this point, the solution is passed back to the planning component

---

[6] An additional reason to omit the modeling of computational errors is that the subjects whose performance we have studied rarely make these errors. All of our subjects have been college students with poor problem-solving skills. They are proficient in computation but not in problem solving.

which then must determine whether the entire problem has been solved or only a sub-problem. A number of things feed into this determination. First, the connectionist network is called upon to find any other plausible situations after the first one has been removed from the problem. A check is carried out to see if there are additional unknowns anywhere in the known problem structure. If potential subproblems are discovered, their clauses and relevant input information is fed back into the model, and the entire cycle begins again. If no additional sub-problems are recognized, the model produces as its answer the computed value for the last sub-problem it solved. Table 1 contains a complete trace of the model's activity for a multi-step problem.

The hybrid model is able to solve problems having more than one unknown. Such problems are common in arithmetic and algebra, and they are frequently studied because students do not routinely solve them easily. The different components of the model pass information back and forth as necessary. For some problems, a re-cycling through the connectionist network will be unnecessary because only one configuration will be possible. For other problems, the model may move back and forth between the connectionist network and the production systems until it develops enough information to create a workable plan.

Thus far, the hybrid model successfully solves problems of the type illustrated in Table 1. Extensions of the model to deal with more complex problems are ongoing, as are comparisons of human and model solutions. The initial findings are encouraging. The hybrid model presented here can solve single or multiple-step problems, and it produces solutions that appear similar to human subjects' solutions.

Table 1: Hybrid Model Output for a Multi-Step Problem

Model Output:	*Annotated Description of Output:*
Joe won $100 in the state lottery. He spent some of it on toys for his two children. He bought a doll for Sue that cost $25 and he bought a stuffed bear for Ellen that cost $28. How much of his lottery winnings did he have after he bought the toys?	Problem Text

```
1 1 1 0 1 0 0 0 0 0 1 0 1 0 0 0 0 1 0 0 1 0 0 0 1 0 0
```
Input Vector for Connectionist Model

owner	Joe person	First Clause
object	dollars dollars 100.000000 none	
time	0	

owner	Joe person	Second Clause
object	dollars dollarsUNKNOWN decrease; toys toys UNKNOWN none	
time	1	

owner	Joe person	Third Clause
object	amount dollars 25.000000 decrease; doll toys 1.000000 increase	
time	1	

owner	Joe person	Fourth Clause
object	amount dollars 28.000000 decrease; stuffed_bear toys 1.000000 increase	
time	1	

owner	Joe person	Fifth Clause
object	dollars dollars UNKNOWN none	
time	2	

`0.369  0.388  0.301  0.280  0.321  ——> GR`

First sub-problem identification by connectionist model.

```
* Combo: 1
  Combo: 2
  Combo: 3
```
The possible configurations. * indicates the one that yields the highest activation value (found via small production system and evaluated with connectionist model).

owner	Joe person
object	dollars dollars UNKNOWN decrease; toys toys UNKNOWN none
time	1

owner	Joe person
object	amount dollars 25.000000 decrease; doll toys 1.000000 increase
time	1

owner	Joe person
object	amount dollars 28.000000 decrease; stuffed_bear toys 1.000000 increase
time	1

The clauses that contribute to the configuration selected by the connectionist model as best. The identification of the GROUP situation and the clause information is passed to the next component of the model which sets the initial goal and determines which values will be used in solving the problem.

Table 1 continued:

Model Output:

Entering ExecuteRules
Production Rule: 25
Goal_Stack:: ID_NUMBER-SUBGROUPS SOLVE
ProblemValues: UNKNOWN 25.00 28.00

Production Rule: 26
Goal_Stack: ID_NUMBER-SUBGROUPS SOLVE
ProblemValues: UNKNOWN 25.00 28.00

Production Rule: 27
Goal_Stack: SOLVE
ProblemValues: UNKNOWN 25.00 28.00

Production Rule: 28
Goal_Stack: ID_PART_GR SOLVE
ProblemValues: UNKNOWN 25.00 28.00

Production Rule: 30
Goal_Stack: SUPERGROUP ID_PART_GROUP SOLVE
ProblemValues: UNKNOWN 25.00 28.00

Production Rule: 32
Goal_Stack: SUPERGROUP ID_PART_GROUP SOLVE
ProblemValues: UNKNOWN 25.00 28.00

Production Rule: 33
Goal_Stack: SUPERGROUP ID_PART_GROUP SOLVE
ProblemValues: 53.00 25.00 28.00

Production Rule: 34
Goal_Stack: ID_PART_GROUP SOLVE
ProblemValues: 53.00 25.00 28.00

Production Rule: 29
Goal_Stack: SOLVE
ProblemValues: 53.00 25.00 28.00

Production Rule: 24
ProblemValues: 53.00 25.00 28.00

Production Rule: 0
Partial Answer = 53.000000

Annotated Description of Output:

IF	{the top goal is SOLVE, the situation is GROUP, and the number of subgroups is not known}
THEN	{add a new goal of identifying the number of subgroups.}
IF	{the number of subgroups is unknown and the goal is to find the number of subgroups}
THEN	{count the number of subgroups and store the value}
IF	{the goal is to find the number of subgroups and that number is now known}
THEN	{delete the goal from the goal stack}
IF	{the goal is SOLVE, the situation is GROUP, the number of subgroups is known, and there is an unknown in the problem}
THEN	{set a new goal to find out which part of the problem is unknown}
IF	{the answer is unknown, the goal is to identify where the unknown is located, and if it is in the supergroup location}
THEN	{add the goal of computing the supergroup to the goal stack}
IF	{the answer is unknown, and the goal is to find the supergroup}
THEN	{add all subgroup values and store the result as the answer}
IF	{the goal is to find the supergroup, and the answer is known}
THEN	{store this information in the problem values}
IF	{the goal is to find the supergroup and it is known}
THEN	{delete the goal from the goal stack}
IF	{the goal is to identify the missing part of a group problem but there are no missing parts}
THEN	{delete the goal from the goal stack}
IF	{the goal is to solve the problem but there are no unknowns}
THEN	{remove the goal from the goal stack}
IF	{the goal stack is empty}
THEN	{return the answer}

The first sub-problem has been solved.

Table 1 continued:

Model Output:

0.481 0.357 0.381 0.370 0.368 ——> CH

* Combo: 1

owner	Joe person
object	dollars dollars 100.000000 none
time	0

owner	Joe person
object	dollars dollars 53.000000 decrease;
	toys toysUNKNOWN none
time	1

owner	Joe person
object	dollars dollars UNKNOWN none
time	2

Entering ExecuteRules
Production Rule: 1
Goal_Stack: ID_PART_CHANGE_SOLVE
ProblemValues: 100.00 53.00 UNKNOWN

Production Rule: 16
Goal_Stack:END ID_PART_CHANGE SOLVE
ProblemValues: 100.00 53.00 UNKNOWN

Production Rule: 18
Goal_Stack: END ID_PART_CHANGE SOLVE
ProblemValues: 100.00 53.00 UNKNOWN

Production Rule: 19
Goal_Stack: END ID_PART_CHANGE SOLVE
ProblemValues: 100.00 53.00 47.00

Production Rule: 20
Goal_Stack: ID_PART_CHANGE SOLVE
ProblemValues: 100.00 53.00 47.00

Production Rule: 5
Goal_Stack: SOLVE
ProblemValues: 100.00 53.00 47.00

Production Rule: 4
ProblemValues: 100.00 53.00 47.00
Partial Answer = 47.000000

Final Answer = 47.000000

Annotated Description of Model Output:

At this point the system re-examines the original input to determine if there are other situations containing other problems to be solved. It finds a CHANGE, and there is only one possible configuration.

The necessary clauses are identified for the planning and execution components.

The system recognizes that there is an unknown value for the object toy but disregards it in favor of the selected change configuration.

The production begins a new cycle:

IF {the goal is to solve the problem; the situation is *CHANGE;* and there is an unknown value on the value_list}

THEN {add a new goal of identifying which part of the *Change* situation is unknown}

IF {the goal is to identify the which part of the problem has an unknown and if the last element of the value_list is unknown}

THEN {add a new goal of finding the end_result}

IF {if the goal is to find the end_result and the direction of change is negative}

THEN {set ANSWER to the difference between the start_amount and the amount_of_change}

IF {the goal is to find the end_result and there is only one unknown value in the value_list and if a value is known for ANSWER}

THEN {replace the unknown in the value_list with the value of ANSWER}

IF {the goal is to find the end_result and there are no unknowns in the value_list}

THEN {delete the goal from the goal stack}

IF {the goal is to identify a missing part but all parts are known}

THEN {delete the goal from the goal stack}

IF {the goal is to solve the problem but there are no unknowns}

THEN {pop the goal from the goal stack}

IF {the goal stack is empty}
THEN {return the answer}

References

Anderson, J. A. (1990). Hybrid computation in cognitive science: Neural networks and symbols. *Applied Cognitive Psychology, 4,* 337-347.

Bereiter, C. (1991). Implications of connectionism for thinking about rules. *Educational Researcher, 20,* 10-16.

Johnson-Laird, P. (1988). *The computer and the mind.* Cambridge, MA: Harvard University Press.

Marshall, Sandra P. (1990). The assessment of schema knowledge for arithmetic story problems: A cognitive science perspective. In G. Kulm (Ed.), *Assessing higher order thinking in mathematics.* Washington, D.C.: AAAS.

Marshall, Sandra P. (in press). Assessing schema knowledge. In N. Frederiksen, R. Mislevy, & I. Bejar (Eds.), *Test Theory for a New Generation of Tests.* Hillsdale, NJ: Lawrence Erlbaum Associates. [a]

Marshall, Sandra P. (in press). Assessment of rational number understanding: A schema-based approach. In T. Carpenter, E. Fennema, & T. Romberg, *Rational Numbers: An Integration of Research.* Hillsdale, NJ: Lawrence Erlbaum Associates. [b]

Marshall, Sandra P. (in press). Statistical and cognitive models of learning through instruction. To appear in Meyrowitz, A. L. & Chipman, S. (Eds.), *Cognitive Models of Complex Learning.* Norwell, MA: Kluwer Academic Publishers. [c]

Marshall, Sandra P., Barthuli, Kathryn E., Brewer, Margaret A., & Rose, Frederic E. (1989). *STORY PROBLEM SOLVER: A schema-based system of instruction.* Technical Rep. CRMSE 89-01 (ONR Contract N00014-85-K-0661), San Diego: San Diego State University.

Marshall, Sandra P. (1991). *Computer-Based Assessment of Schema Knowledge in a Flexible Problem-Solving Environment.* Technical Rep. CRMSE 91-01 (ONR Grant N00014-89-J-1143), San Diego: San Diego State University.

McClelland, J. L., & Rumelhart, D. E. (1989). *Explorations in parallel distributed processing: A handbook of models, programs, and exercises.* Cambridge, MA: The MIT Press.

Minsky, M. (1991). Logical versus analogical or symbolic versus connectionist or neat versus scruffy. *AI Magazine, Summer 1991,* 35-51.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In Rumelhart, D. E. & McClelland, J. L. (Eds.), *Parallel distributed processing, Volume 1: Foundations* (pp. 318-362). Cambridge, MA: The MIT Press.

Schneider, W. & Oliver, W. L. (in press). An instructable connectionist/control architecture: Using rule-based instructions to accomplish connectionist learning in a human time scale. To appear in K. VanLehn (Ed.), *Architectures for intelligence.* Hillsdale, NJ: Erlbaum Associates.

Winger, E. (1987). *Artificial intelligence and tutoring systems.* Los Altos, CA: Morgan Kaufmann Publishers, Inc.

Dr. Richard Lesseh
ACT
P.O. Box 168
Iowa City, IA 52243

Dr. Bill F. Lehers
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213-3890

Dr. Stephen Kosslyn
Harvard University
1236 William James Hall
33 Kirkland St.
Cambridge, MA 02138

Dr. Jeffrey L. Kennington
Computing Science & Engr.
School of Engr. & Applied Science
Southern Methodist University
Dallas, TX 75275

Dr. John Jonides
Department of Psychology
University of Michigan
Ann Arbor, MI 48104

LTER/C D. F. Lyon
Chief of Naval Applied
Research Outfit
Naval Command Headquarters
FMO Reffin, N.J. BK 2DD
CANADA

Dr. Paul E. Lehner
Department of Information
Systems & Engineering
George Mason University
4400 University Drive
Fairfax, VA 22030-4444

Dr. Renata Kotovsky
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

Dr. Henry Khrehanten
Naval Institute of
Mental Health, D/BRB
5600 Fishers Lane
Parkham Building,
Rockville, MD 20857

Dr. David MacGregor
Decision Research
1201 Oak St.
Eugene, OR 97401

Dr. Richard Lark
Educational Testing Service
Princeton, NJ 08541

Dr. Richard J. Koubek
School of Indus. Engr.
Grissom Hall
Purdue University
West Lafayette, IN 47907

Dr. Giorgio B. Moretty
Department of Management
Science & Information
College of Education
University of Maryland
College Park, MD 20742

Dr. Patrick Kyllonen
AFHRL/MOEL
Brooks AFB, TX 78235

Dr. David Klann
Tinhaled Cognition Program
TIDAL Bldg., 2900 Braddock Blvd.
University of Michigan
Ann Arbor, MI 48109-2108

Dr. Jeremy Kilpatrick
Dept. of Mathematics Education
105 Aderhold Hall
University of Georgia
Athens, GA 30602

Dr. Sung-Ho Kim
Educational Testing Service
Princeton, NJ 08541

Library
Naval Training Systems Center
1290 Research Parkway
Orlando, FL 32826-3224

Dr. Herbert Liu
Naval Academy of Science
2101 Constitution Ave
Room NA-169
Washington, DC 20418

Van M. Mehn
NPRDC Code 143
San Diego, CA 92152-6800

Dr. Ann Meier
Mail Code ER2
NASA Ames Space Center
Houston, TX 77058

Dr. Christen Lisso
BL
2550 Hanover Street
Palo Alto, CA 94304

Dr. Marcie C. Lin
Graduate School
of Education, EMST
Tolman Hall
University of California
Berkeley, CA 94720

Dr. William L. Maloy
NETPMSA, Code 04
Pensacola, FL 32509-5000

Dr. Bruce Meltzer
George Mason University
4400 University Drive
Fairfax, VA 22030

Dr. John Juola
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

Dr. Marcel Just
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. D. Gage Kingsbury
Portland Public Schools
Research & Evaluation Dept.
501 North Dixon Street
P.O. Box 3107
Portland, OR 97209-3107

Dr. M. Diane Langston
ICL Nele Austin
11490 Commerce Park Drive
Reston, VA 22091

Dr. Mary Lamon
University of North Carolina
Dept. of Computer Science
CB #3175
Chapel Hill, NC 27599

Dr. Mary Mattin
Directorate, Instr. Tech.
HQ USAF/ADTET
USAF Academy, CO 80840-5017

Dr. Chantel J. Mania
Head, DoD Coordinator,
Startfing Personnel Program
Stasis Code: PERS 2WF234
Navy Annex, Room 2832
Washington, DC 20350

Dr. Hendrik Mahn
ALHRA, Bay 44
Vinton AFB
AZ 85240

Dr. Rob Keehle
University of Minnesota
Department of Psychology
Elliot Hall
75 E. River Road
Minneapolis, MN 55455

Dr. Michael Kaplan
Office of Basic Research
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333-5600

Dr. Steven S. Klutchman
Code 2212, Building 1171/1
Naval Underwater Systems Center
Newport, RI 02841

Dr. Kenneth A. Kirvington
The Salk Institute
P.O. Box 85800
San Diego, CA 92186-5800

Logicon Inc. (Attn: Library)
Tactical & Training Systems Div.
P.O. Box 85158
San Diego, CA 92138-5158

Prof. David F. Lohman
College of Education
University of Iowa
Iowa City, IA 52242

Matt Dr. Christian Lahrmann
Kompetenzzentrum Direkt
Hauk Thonfeldr-Kerasen
1130 Wien
AUSTRIA

Dr. Martin J. Ippel
Center for the Study of
Education and Instruction
Leiden University
P.O. Box 9555
2300 RB Leiden
THE NETHERLANDS

Dr. Janne Andrew
Decision Science
Consortium, Inc.
Glen Rimburst St.
97035 Garinberg
THE NETHERLANDS

Z. Andrew
Board of Assessment Consultants
766-388 Larche Ave., W.
Groer/Chests H3N 2W3
CANADA

Dr. Peter Janssen
Res. and Computing Lab. Dept.
Bio. and Computer Sci. Center
University of South Carolina
Columbia, SC 29208

Dr. Gisela Redes
University de Quebec
Montreal, C/CAR
Montreal, Quebec H3C 3P8
CANADA

Dr. Robbie Makaba
Naval Postgraduate School
Naval Postgraduate School
Monterey, CA 93943

Dr. Douglas Kelly
University of North Carolina
Department of
Statistics, CB #3260
Chapel Hill, NC 27599

Dr. Axel L. Koenker
Georgia Institute of Technology
College of Computing
Atlanta, GA 90332-0280

Dr. Yeh-Jung Lee
Department of Computer Science
Code CS/LE
Naval Postgraduate School
Monterey, CA 93943

Dr. Sylvia Koreban
University of Michigan
Mental Health Research Institute
205 Washtenaw Place
Ann Arbor, MI 48109

Dr. A. Konstantakelo
NRC-CDU
17 Charles Street
London
ENGLAND WC1H OAH

Dr. Marie J. Koa
PIC 302 Box 13
PPO AE 09499-1500

Dr. Steven W. Keele
Department of Psychology
University of Oregon
Eugene, OR 97403

Dr. Wesley Kollwy
IBM T. J. Watson
Research Ct.
P.O. Box 704
Yorktown Heights, NY 10598

Dr. J.A.S. Keke
Center for Computer Systems
Building MT1
Florida Atlantic University
Boca Raton, FL 33431

Dr. Daniel B. Ammo
US Nuclear Regulatory
Commission
NMRR 108
Washington, DC 20555